INSTITUTE OF PHYSICS PUBLISHING PHYSICS IN MEDICINE AND BIOLOGY

Phys. Med. Biol. **51** (2006) 3941–3952 doi:10.1088/0031-9155/51/16/004

A modified simulated annealing algorithm for parameter determination for a hybrid virtual model

C Carletti[1], **P Meoli**[2] **and W R Cravero**[1]

[1] Departamento de Física, Universidad Nacional del Sur, Av. Alem 1253, Bahía Blanca, Argentina
[2] Instituto Alexander Fleming, Buenos Aires, Argentina

E-mail: wcravero@uns.edu.ar

Received 6 January 2005, in final form 14 June 2006
Published 26 July 2006
Online at stacks.iop.org/PMB/51/3941

Abstract

In this work we present an algorithm based on the simulated annealing (SA) method for electron beam spectrum reconstruction from central axis PDD data. We use a simulated beam in order to assess the accuracy of the method, and compare it with others usually employed. We found that our modified SA algorithm produced excellent reconstructed PDDs and beam profiles and improved reconstructed spectra. We also tested our method for the simultaneous determination of an electron energy spectrum and the position of a trimmer from central axis PDD data, and found this to be sufficient for the determination of the physical parameters of this hybrid model.

1. Introduction

The Monte Carlo method is increasingly being used as an accurate and practical approach for electron and photon beam dose calculation, and clinical use is beginning to gain ground (Andreo 1991, Ma and Jiang 1999).

It is well known that treatment of electron transport in an inhomogeneous medium can lead to large discrepancies in dose distributions between conventional analytical algorithms and measurements (see for example Ma *et al* (1999)). This makes Monte Carlo methods particularly appealing for calculations for clinical electron beams.

When using the Monte Carlo method for dose calculations, it is necessary to generate huge phase-space data in order to take into account the physical properties of the primary beam as well as the effects of the accelerator treatment head within it. From a practical point of view, relevant experience in Monte Carlo calculations is generally required in order to perform reliable MC calculations. This level of skill is not readily available in most radiotherapy institutions. Even provided the required expertise is available, it is very time consuming to generate a complete useful set of phase-space data for clinical use.

The energy distribution of the initial electron beam depends in general on the linear accelerator type, the method of electron injection, the tuning of the accelerator tube and the beam handling system (Deasy *et al* 1996). Furthermore, in order to generate the phase-space data, accurate and reliable knowledge of both the primary beam energy distribution as well as for treatment head geometry and materials are necessary (Deasy *et al* 1994, Rogers 1991).

In order to overcome these problems, virtual source models have been developed. In these models, the treatment head is basically replaced by a small number of virtual sources and simple structures which, appropriately placed, will generate phantom surface phase-space data from which Monte Carlo calculations of absorbed dose can proceed, and which ultimately generate dose distributions close enough (within prescribed tolerances) to those yielded by the whole treatment head (Faddegon and Blevis 2000, Deng *et al* 2001).

The spatial configuration and the energy spectrum are to be determined for each of the virtual sources. One way to do that is simply to compare calculated doses from the virtual source results with complete Monte Carlo calculations that incorporate the actual treatment head. Another way to achieve this goal would be to experimentally determine virtual sources' energy spectra from dosimetric measurements from a water phantom in a clinical accelerator. Experimental determination of energy spectra are not usually performed on therapeutic electron beams, but a few measurements with dedicated magnetic spectrometers have been reported (Deasy *et al* 1994, 1996, Kok and Welleweerd 1999). In these cases the actual measurement took place at some point after the beam had passed through the structures of the treatment head.

Several methods have been proposed to derive the energy spectra both for photon and electron beams from measured PDDs (Deng *et al* 2001, Faddegon and Blevis 2000, and references therein).

While reproduced PDDs show a very good agreement with measurements, the match in energy spectra against independently measured spectra or those calculated using Monte Carlo and incorporating the actual treatment head, is not so good. Issues that must be addressed before using a virtual source model are the optimization methods that are to be used in order to define the physical characteristics of the sources (energy spectrum, position, shape in the case of extended sources, etc), and the set of measurements necessary to determine the model (Siljamki *et al* 2005). Several prescriptions have been published on how to do this (see for example Jiang *et al* (2000)) in the framework of Monte Carlo TPSs commissioning.

A related approach (which we will call the *hybrid model*) is to replace the entire accelerator head with a simpler combination of virtual sources and structures through which radiation from the virtual sources passes. This method can be seen as a compromise between a pure virtual source model for the head (which might be tricky to define) and the simulation of the entire treatment head. In particular it can be used when electron applicators are incorporated into treatment heads.

Once the method of simulation is chosen, the problem reduces to one of optimization, where an objective function (OF) that, in general, represents the difference between measured and calculated doses, is minimized.

One question that arises when multiple sources and structures are used to simulate the head is whether there is more than one configuration that yields the same minimum value (within statistical uncertainty) for the OF. This is not easy to answer since it depends on which set of measurements is considered. For a sufficiently large set of measurements there should be only one configuration of virtual sources that reproduces all of them. Recently proposed virtual source models have considered central axis PDDs together with horizontal and oblique measurements in order to define virtual source models (Siljamki *et al* 2005). From a practical

point of view, it is of course desirable to keep the set of measurements necessary to determine the source model to a minimum.

In this work we will study the use of central axis PDDs in order to determine a simple hybrid model and to see whether the model can be uniquely determined from central axis PDDs only. We will use a simulated annealing method (SA) which is well suited for the optimization of systems exhibiting a complex landscape for their OF.

Our aim is to assess the accuracy of the employed SA method and therefore determine whether it is possible to use it together with central axis PDDs in finding the physical parameters of a simple hybrid model. However, instead of simulating an actual treatment head, we will use as our starting point calculated dose data for the same hybrid model using PENELOPE. In this way, we know *a priori* the target spectrum, so we can study the reconstruction of PDDs and the reconstruction of target spectra. Besides, experimental errors in measuring energy spectra are eliminated, and statistical fluctuations in the Monte Carlo simulation of an entire treatment head are much reduced. We test the optimization method in an ideal case where it should be able to find the parameters for the system exactly. Discrepancies can then be blamed on the method and will show if our hypothesis—that the system can be determined from PDD data alone—can be tested further on a more realistic system.

The paper is organized as follows. In the following section we recall the SA method and define the system we use to test it. In section 3 we will test the SA method by reproducing the electron energy spectrum of a simple electron source, starting from the known central axis PDD data in a water phantom. Its intrinsic accuracy will be compared to different methods used to find energy spectra. We then apply the SA method to establish a simple hybrid model and determine whether such a system can be derived from PDDs measurements.

2. Methods and materials

As we have already said, we will use a simple theoretical model in order to test our optimization method. The system consists of a point electron source with a given angular aperture θ, whose energy spectrum is known *a priori*. Its SSD is set to 100 cm and an Al cylindrical trimmer of given aperture and width is placed between the source and the water phantom. We will choose small apertures for the trimmer. That condition defines a worse case configuration. Most of the contribution to central axis PDD is in this case primary radiation. Our optimization method should be able to guess the correct source spectrum together with the position of the trimmer from central axis PDD data. However, in the limiting case of a null aperture, i.e., primary radiation only, the central axis PDD would become independent of the trimmer position. On the other hand, a large trimmer aperture implies a larger contribution to central axis PDDs from scattered radiation. In general, the contribution of the scattered radiation from the phantom volume to the central axis PDD would be (for a given aperture) a function of the trimmer position, thus making it a lot easier for the algorithm to find out.

As is well known, as a consequence of beam interactions (absorption and bremsstrahlung) with the materials of the accelerator head (primary and secondary collimators, filters, monitor chambers, etc), any real clinical electron beam will be accompanied by a photon field that contributes to the absorbed dose in the water phantom. At this stage we are interested in benchmarking our inverse calculations methods, so we have not included a photon virtual source. However, photons generated by bremsstrahlung in the phantom are of course present and contribute to the PDD. Unlike other methods (see for example Deng *et al* (2000)) we do not need to subtract the photon tail in order to find the electron source spectra.

Our optimization code is based upon the simulated annealing (SA) method (Press *et al* 1986). SA is a well-known technique suitable for optimization of large scale problems, where

the global extreme might be hidden among many local extremes. It is specially useful in cases where the OF to be minimized/maximized is a discrete, very large, configurational space. In our case, we have a multiobjective problem to solve. Our goal is to find a set of best values for the parameters that define the electron energy spectrum and the trimmer position, in order to obtain a satisfactory reconstruction for the target PDD. In general, the information we would have about the real system is the measured PDD (target PDD). In our case, this target PDD is a Monte Carlo calculated PDD. Therefore, the OF should be dependent on how well the SA-reconstructed PDD matches to the target PDD. Most non stochastic optimization methods converge to a quick nearby solution, going downhill as fast as they can. These procedures will lead to a local extreme, but not necessarily to the global one. SA, on the other hand, is an algorithm suitable for finding the optimal configuration by being able to get out of a local extreme.

The SA method is based in the thermodynamic process by which liquids crystallize or metals cool and freeze into a crystallographic structure corresponding to the minimum energy. At high temperature the atoms in a liquid have many allowed thermodynamic states but as the temperature falls the thermal mobility is gradually lost. If the cooling process is carried out slowly enough, the atoms will form a stable structure. In contrast, a rapid cooling rate will lead to a state with many defects and irregularities in the metal structure (this situation usually corresponds to a local energy minimum, i.e., a metastable state). Conventional minimization algorithms may be compared to the rapid cooling or quenching, since in all these cases the algorithms converge to the nearest local minimum.

In any SA algorithm the temperature is a defined factor that determines the randomization degree in the variable parameters, that is, as the temperature falls off the probability of uphill changes in the OF value decreases. Even though the goal of the algorithm is to minimize the OF, in the simulated annealing process not all the random steps are taken in that direction, sometimes random steps up are allowed. While the algorithm always accepts a change that lowers the OF (downhill steps), steps in the opposite direction are only taken with a given probability which depends on the system *temperature* and the magnitude of the configurational change.

Pursuing the analogy with the physical cooling process, the probability that in the SA process a change from the energy E_1 to energy E_2 of the configurational states takes place is given by the function $p = \exp\left(-\frac{(E_2-E_1)}{KT}\right)$, where k is the analogue to Boltzmann's constant, T represents the system temperature and E is the OF value. If $E_2 > E_1$, (i.e., the OF increases), the step will sometimes be taken, depending on the value of p. The routine that decides whether to accept a reconfiguration is based on the Metropolis algorithm, which lies at the heart of any SA algorithm. If $\Delta E > 0$ Metropolis' return is true only with probability $\exp\left(-\frac{\Delta E}{KT}\right)$. If $\Delta E < 0$, the Metropolis result is always true, i.e., if $E_2 < E_1$, the OF decreases and the system will always accept this step. In our case the minimization is preceded by a randomization process over a set of discrete positions available to the trimmer and the weighting factors of the set of monoenergetic electron beams that build up the electron energy spectrum. The OF will be related to the quadratic departure of the trial central axis PDD SA calculated central axis PDD with respect to the real (target MC calculated central axis) PDD. In figure 1 a schematic flow chart of the algorithm structure is shown.

Input data for our algorithm are the MC calculated central axis PDD data file and a PDD data base corresponding to n equally spaced monoenergetic electron beams covering the relevant energy range, for each available position of the trimmer. In our case, both the original MC calculated PDDs as well as the monoenergetic PDDs were calculated using the PENELOPE Monte Carlo code (Salvat *et al* 2001). The geometrical configuration of our system is shown in figure 2.

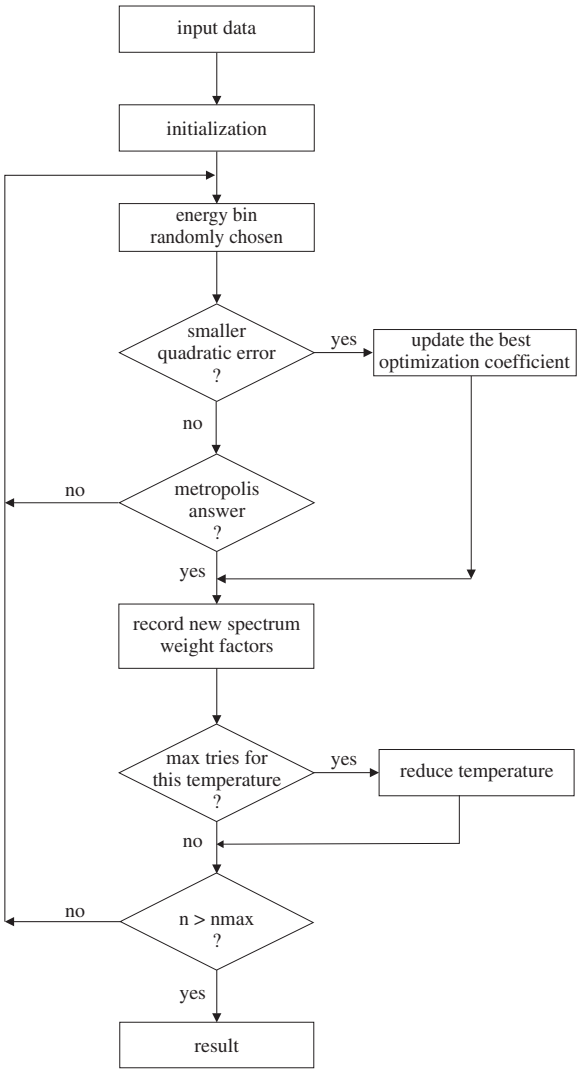

Figure 1. Flow chart scheme for the simulated annealing (SA) algorithm used.

During initialization, starting trial weights for the electron beam spectrum are established as well as the percentage by which the energy of a randomly chosen energy bin is changed. To begin the annealing process, one of the energy bins is randomly selected. Next, its weighting factor is increased or decreased by a small percentage. Central axis PDD is then recalculated from the monoenergetic PDD data. In order to do that, the dose at a certain depth z_i in the water phantom is simply expressed as

$$D'(z_i) = \sum_{j=1}^{n} W_j \cdot D_j^{\text{mono}}(z_i), \tag{1}$$

where $D'(z_i)$ represents the calculated dose at z_i, W_j is the weighting factor corresponding to the jth energy bin and $D_j^{\text{mono}}(z_i)$ is the dose at z_i due to the corresponding monoenergetic

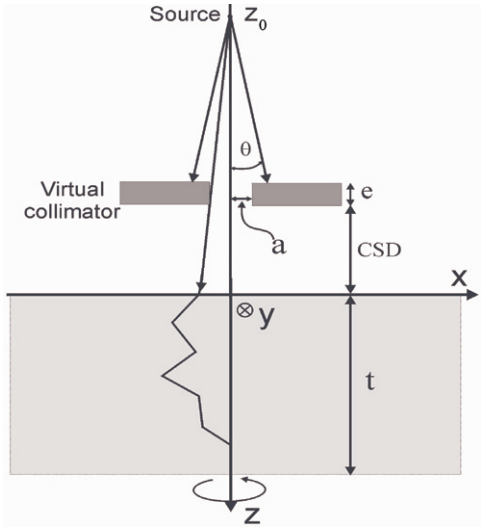

Figure 2. Geometry used to test the spectrum reconstruction.

beam. Every time a change in the spectrum is performed, the OF value is calculated and compared with the last result obtained so far. Our OF is a measure of the departure of the trial PDD from the target PDD, but there is not a unique choice for the OF; Deng *et al* (2001) have defined a correlation coefficient (CC) for the trial PDD, $D'(z)$, and the original PDD, $D(z_i)$:

$$CC = \frac{\sum_{i=1}^{n} D(z_i)D'(z_i)}{\sum_{i=1}^{n} D(z_i)\sum_{i=1}^{n} D'(z_i)}. \tag{2}$$

In our case we have simply used the *quadratic difference*:

$$OF = \sum_{i=1}^{n} (D(z_i) - D'(z_i))^2. \tag{3}$$

We also tested our optimization algorithm using CC as the OF and have found no noticeable difference. We also used $\log(D)$ and $\log(D')$ instead of D and D' in equation (3) but we obtained also very similar results. In order to compare the last PDD data to the previous one, we define the parameter $d_{\mathrm{OF}} = OF_{\mathrm{new}} - OF_{\mathrm{old}}$. If d_{OF} is negative, then the new configuration is better than the previous one, and the spectrum and the OF are updated. On the other hand, if d_{OF} is positive the result is accepted or rejected with a probability $P = \exp\left(-\frac{|d_{\mathrm{OF}}|}{T}\right)$, where the factor T represents the system temperature. T is held constant for a given number of steps. This allows the algorithm to explore the dominion of OF . Then T is decreased, and a new set of steps are taken. As the temperature decreases the probability of accepting a non-favourable change in configuration decreases too. Temperature in an SA calculation is lowered following a given path. In our case it is reduced by a fixed amount after a maximum number of steps have been taken at a given temperature. How steep the temperature descent should be largely depends on the convergence rate of a given optimization problem. We found that a key issue in our annealing algorithm is to allow for slow cooling, although this fact can lead to longer computation times. A reasonable compromise cooling rate is, however, easily achievable. The code exits when a maximum number of iterations have been performed or when single step improvements in the configuration become smaller than a pre-established value.

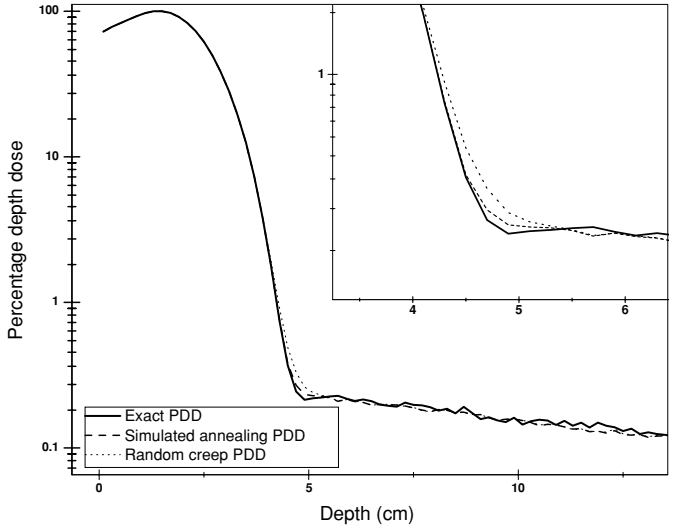

Figure 3. Target and reconstructed PDDs for a point electron source. The source spectrum is shown in figure 4.

Adding a trimmer to the system amounts to an increase in the dimensions of the OF landscape, thus complicating the optimization method. After a short period of instability, the OF corresponding to the correct position clearly becomes smaller than those corresponding to other positions. At an intermediate point we freeze the position whose OF is smaller and then further refine the source energy spectrum as usual. The drawback of this strategy is that it does not scale well if more structures are to be added to the system, which is a logical step if we want to simulate an actual applicator, for example. However, relative distances between a given applicator's structures are supposed to be fixed, so even when a simulated applicator might not have the same shape and size as the actual one, the problem is still reduced to the addition of one extra dimension.

3. Results and discussion

As has been said, our method is checked by using it to find a known original spectrum. We first propose a broad energy beam distribution. Our set of monoenergetic beams used by the algorithm is composed of $n = 20$ monoenergetic beams covering an electron energy range from 0.5 to 10 MeV.

The initial trial spectrum is in our case constructed from the original PDD by calculating its R_{50}. Its most probable spectral energy E_p is determined from the well-known relation,

$$E_p = 2.5224R_{50} - 0.252. \tag{4}$$

A unity weighting factor is assigned to the E_P energy bin, while for the rest of the energy bins the weighting factors were distributed as follows: 0.5 for the two E_p adjacent energy bins, 0.25 for the next neighbours, and 0.01 for the remaining energy bins.

In figure 3 we show PDD data from the spectrum reconstructed by using SA and the random creep (RC) algorithm of Deng *et al* (2001), and we compare them with the original beam's PDD data. The RC method is quite similar to our algorithm except that it only accepts downhill steps. There is no temperature parameter, or in other words it is a zero temperature

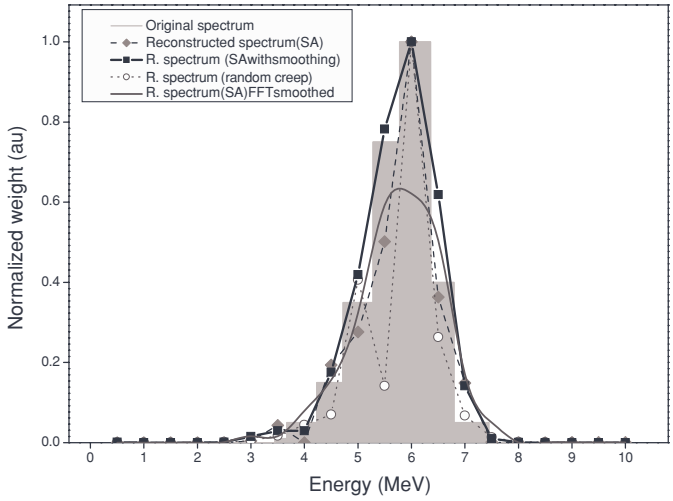

Figure 4. Target and reconstructed spectrum for a narrow spectrum 6 MeV electron beam, using different optimization methods (see text for details).

algorithm, and in fact we could reproduce its results by simply making $T = 0$ in the SA algorithm. Both algorithms generally agree very well with the original beam's PDD data. However, we note a marginally better agreement using SA, particularly in the region where the bremsstrahlung tail begins to be important. As has been said, in our system we only have bremsstrahlung originated by the phantom, as we are considering a pure electron source. Consequently, we have not subtracted (in any of the methods) the bremsstrahlung tail from the original beams data in order to perform the spectrum determination.

In figure 4 we see the corresponding spectrum reconstructed using SA and RC, as well as the original target spectrum (shown in the figure as a shadowed histogram for the sake of clarity). We note that discrepancies with the target spectrum here can only be due to problems with the optimization algorithms or due to the fact that we use a finite number of monochromatic beams to simulate a continuum. In our case both algorithms use the same number of monochromatic beams, so we note that SA seems to be more accurate than RC. RC appears to fall within local minima that correspond to noisier spectra from which it cannot escape because it cannot take uphill steps. We also note that the fitting is not perfect using SA either. We believe the origin of the differences is related to the flatness of the OF minimum, i.e., there are several configurations that yields almost the same value for the OF within the computational tolerance, although they may represent quite different spectra. That might not be detectable at PDD level, since the quadratic difference in PDD *is* the OF. Nevertheless one might be interested in the spectra in addition to PDDs, particularly if the reconstructed spectrum is to be used as the starting point of MC dosimetric calculations. We noticed that some degree of noise in the reconstructed spectrum can be eliminated by slightly modifying the OF, incorporating the idea of smoothness in the SA algorithm. A spectrum with a sharp discontinuity should be slightly penalized (in terms of the OF value) compared to a smooth one if both generate the same PDD within tolerances. In this way actual spectrum discontinuities can be taken into account but those originated in the flat structure of the OF minimum should be reduced. There are several different ways to achieve this. We simply chose to add to any energy bin weight a bit of its nearest neighbours in each SA iteration. For comparison, we also

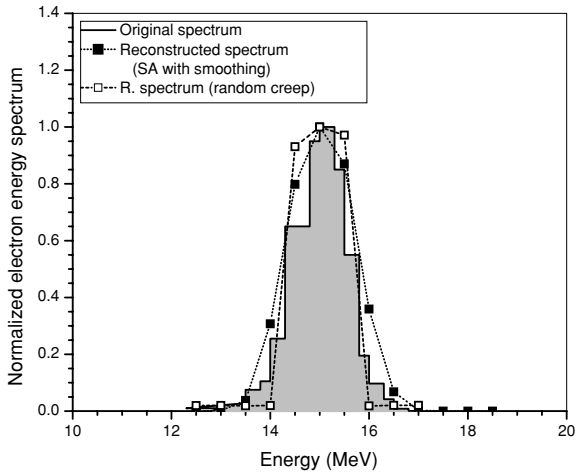

Figure 5. Same as figure 4 for a narrow spectrum 15 MeV electron beam.

show in figure 4 the result of applying a standard FFT based smoothing algorithm *after* the spectrum has been determined using SA. We see that our method is still capable of obtaining the sharp peak in the spectrum maximum, which standard methods cannot.

We find that when beam energy is increased, reconstruction of the spectrum is faster and more efficient. In figure 5 we show spectrum reconstruction for a higher (15 MeV) beam energy. We see that both SA and RC reconstruction algorithms give reasonable results, although smoothed SA still yields the best results.

In order to quantify the accuracy of the spectrum reconstruction let us introduce a measure of the departure from the target spectrum as given by SA and RC, and what we call smoothed simulated annealing (SSA). We define the following percentage relative error e:

$$e = \frac{\int |W_{\text{reconstructed}}(E) - W_{\text{target}}(E)|\,\mathrm{d}E}{\int W_{\text{target}}(E)\,\mathrm{d}E} \times 100. \tag{5}$$

For SA we get $e \cong 27\%$, while for RC $e \cong 36\%$ and SSA yields $e \cong 8\%$. These may look like rather large discrepancies for SA and RC, taking into account that a perfect fitting could be expected in this case, since the target spectrum is a simple theoretical one. However, If we look for R_{50} in central axis PDDs, we find that reconstructed R_{50} lies within 0.01 mm from target R_{50}. In data reported by Deng *et al* (2001), relative error for the reconstructed spectrum are around 50%, and still R_{50} values are within 0.1 mm from their target PDD R_{50}.

We find that when beam energy is increased, reconstruction of the spectrum is somewhat less efficient. In figure 5 we show spectrum reconstruction for a 15 MeV beam. We see that both SA and random creep reconstruction algorithms give reasonable results, although Smoothed SA still yields the best results. The relative errors obtained for SSA are $e \cong 32\%$, while RC gives $e \cong 41\%$. We see that SSA results, even when poorer than those obtained for 6 MeV, are still better than those yielded by our RC calculations and those reported by Deng *et al* (2001).

In figure 6 we show the optimization process for different positions of an Al trimmer. We see that, after a relatively small number of iterations, the OF for the correct position of the trimmer can clearly be distinguished as the lower one. That means that central axis PDD data is in principle sufficient to determine the position of a trimmer, without resorting to a

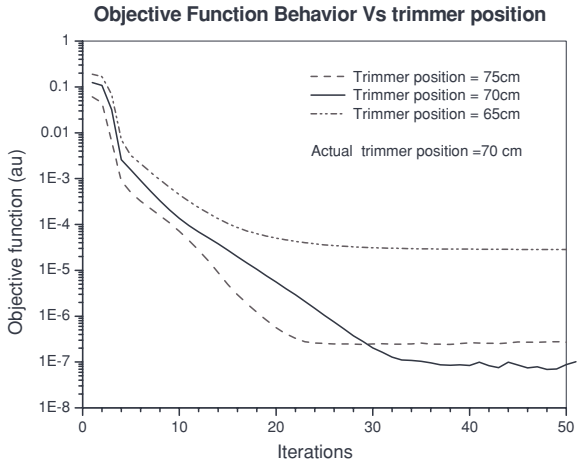

Figure 6. Evolution of the objective function (OF) for the optimization of the spectrum and trimmer position as a function of the number of iterations at different temperatures.

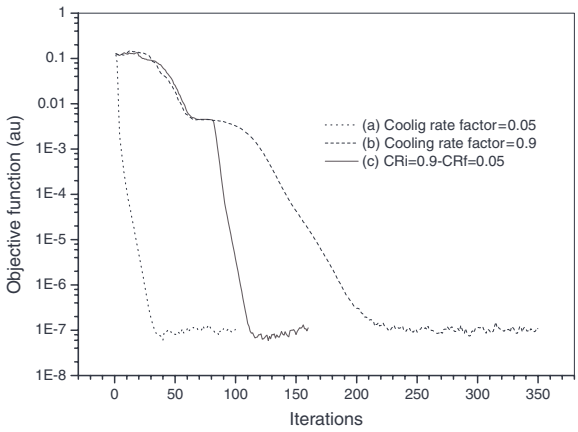

Figure 7. Same as figure 5 for a given trimmer position for different cooling rates.

more complicated multi-objective scheme. We observe that the evolution of the OF is not monotonic but exhibits a wavy behaviour as it approaches its minimum.

Depending on the cooling rate, the OF's evolution as the number of iterations increases can even show a distinctive plateau before falling to its minimum value, as figure 7 shows. This plateau could be due to a local minimum in the OF, from which it is known that SA methods can get out but other non-stochastic methods could get trapped, so it deserve further exploration. The bigger the cooling rate, the less stochastic the SA method becomes. We see that the plateau vanishes as the cooling rate increases. That means that the plateau is not a local minimum but a region with a gentler energy landscape. We further confirm this by changing the temperature cooling rate at the moment the OF reaches the plateau (solid line in figure 7). We see that the OF leaves the plateau immediately after the cooling rate change. Trying to identify the physical origin of the plateau, we parameterized the spectrum using the position of the maximum, the

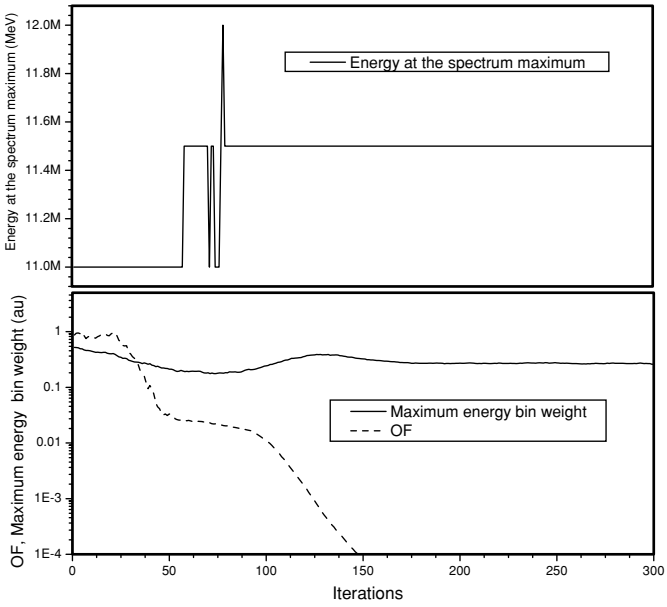

Figure 8. Evolution of the OF, position of the maximum spectrum energy, and its bin weight when the OF reaches a plateau.

value of the weight at the maximum, and the spectrum FWHM. In figure 8 we show the evolution of these parameters together with the OF value as the SSA algorithm evolves. We see that the existence of the plateau can be related in this case to a lack of definition in the position of the spectral maximum in the spectrum, which oscillates when the OF reaches the plateau. As the optimization evolves, the position of the maximum gets fixed, and the OF slope increases as the algorithm continues tuning the rest of the parameters.

4. Conclusions

We have used a modified SA method to determine energy spectra for simple electron beams from central axis PDD curves. We have shown that reconstructed source PDD curves show excellent agreement with Monte Carlo calculated PDDs from the original source. Fitted spectra have also been compared with target spectra, and we found that even when the original source is a very simple theoretical one, usual methods are not able to accurately reconstruct the corresponding energy spectrum. We believe this limitation is mainly due to the monochromatic nature of the beams used to reconstruct the target continuous spectrum. However, modified simulated annealing methods such as SSA can greatly improve the results obtained at the expense of making a few additional assumptions about the system. PDD reconstruction on the other hand is very good in all methods tested. This is not a surprise since the OF is related to PDD so its reconstruction is more or less guaranteed, at least for a theoretical source. We tested our method in a hybrid model in order to simultaneously determine the source energy spectrum and the position of a trimmer from central axis PDD data. We showed that the optimization method is able to correctly determine the position of the trimmer. This opens the way for using hybrid models as an intermediate approach between virtual sources and entire treatment head simulation. Further work includes using continuous spectrum beams to

reconstruct energy spectra, adaptive energy bin widths and using SSA methods to determine more realistic hybrid models.

Acknowledgment

This work has been partially supported by Universidad Nacional del Sur, Argentina, under PGI 24/034 and CONICET, PIP 5808.

References

Andreo P 1991 *Phys. Med. Biol.* **36** 861
Deasy J *et al* 1994 *Med. Phys.* **21** 1369
Deasy J O, Almond P R and McEllistrem M T 1996 *Med. Phys.* **23** 675
Deng J *et al* 2000 *Phys. Med. Biol.* **45** 411
Deng J *et al* 2001 *Phys. Med. Biol.* **46** 1429
Faddegon B A and Blevis I 2000 *Med. Phys.* **27** 514
Jiang S B, Kapur A and Ma C-M 2000 *Med. Phys.* **27** 180
Kok J G M and Welleweerd J 1999 *Med. Phys.* **26** 2589
Ma C-M, Mok E C, Kapur A, Pawlicki T, Findley D O, Brain S, Forster K and Boyer A L 1999 *Med. Phys.* **26** 2133
Ma C-M and Jiang S B 1999 *Phys. Med. Biol.* **44** R157
Press W *et al* 1986 *Numerical Recipes in Fortran* (Cambridge: Cambridge University Press)
Rogers D W 1991 *Int. J. Appl. Radiat. Isot.* **42** 965
Salvat F *et al* 2001 *PENELOPE, A Code System for Monte Carlo Simulation of Electron and Photon Transport* (Paris: OECD Publications)
Siljamki S, Tillikainen L, Helminen H and Pyyry J 2005 *Med. Phys.* **32** 2113